

Chapter 3: Using Enstore

3.1 Copying Files with Encp

Encp is an end-user command used to copy data files from disk to storage media and vice-versa. It is maintained in KITS as a separate product from Enstore, and is designed to be used in conjunction with it. Encp can only be used from on-site machines in the fnal.gov domain. We assume you have UPS/UPD running on your local machine.

3.1.1 Encp Command Syntax and Usage

Encp is very similar to the **cp** command in UNIX, and the **cp** command's behavior has been duplicated in **encp** to a large extent.

Syntax

The syntax is:

```
% encp [<options>] <source_file> <destination_file>
```

With the exception of the option **--help**, we defer the list and definitions of options to section 3.2 *Encp Command Options*, and instead proceed with usage information. Use the **--help** option to request the option listing for **encp**:

```
% encp --help
```

```
USAGE: encp [opts] inputfilename outputfilename
or
encp [opts] inputfilename1 ... inputfilenameN\
outputdirectory
--age-time=
--config-host=
--config-port=
--data-access-layer
--delayed-dismount=
--delpri=
--ephemeral
--file-family=
```

```
--help
--no-crc
--pnfs-is-automounted
--priority=
--usage-line
--verbose=
```

3.1.2 Setup encp

To setup **encp**, run the command:

```
% setup -q <qualifier> encp
```

where **<qualifier>** stands for one of the Enstore systems. Currently, these include:

stken	for general Fermilab users
d0en	for D0 users
cdfen	for CDF users

For example, a CDF experimenter would type:

```
$ setup -q cdfen encp
```

3.1.3 Upload Files to Enstore Media

First, setup **encp**. To copy a file from disk to Enstore, use **encp** to write to the **/pnfs** namespace. You can use filename expansion (wildcard characters to specify a group of files). Run the command as follows:

```
% encp [<options>] /path/to/<localfilename> \
/pnfs/storage-group/dir/in/which/to/copy/<remotefilename>
```

Examples

- 1) Standard copy; no options. Copy **myfile** to the directory **/pnfs/expt1/subdir/**:

```
% encp /path/to/myfile /pnfs/expt1/subdir/
```
- 2) Request a delayed dismount (**--delayed-dismount**) of 5 minutes, and have the process output some information to screen (**--verbose**). Again, copy **myfile** to the directory **/pnfs/expt1/subdir/**:

```
% encp --delayed-dismount=5 --verbose=3 \
/path/to/myfile /pnfs/expt1/subdir/
```
- 3) Copy all the files in the cwd starting with the string **trigger1** to the directory **/pnfs/expt1/subdir/**:

```
% encp trigger1* /pnfs/expt1/subdir/
```

3.1.4 Download Files from Enstore Media

Run encp

First, setup **encp** (with qualifier; see section 3.1.2 *Setup encp*).

To download a file from Enstore (i.e., from storage volume) to disk, use **encp** to read from the `/pnfs` namespace. You can use filename expansion (wildcard characters to specify a group of files). Run the command as follows:

```
% encp [<options>] \  
/pnfs/storage-group/dir/from/which/to/copy/<remotefilename> \  
/path/to/<localfilename>
```

Examples

- 1) Standard copy; no options. Copy `/pnfs/expt1/subdir/myfile` to the cwd, keeping the filename the same:

```
% encp /pnfs/expt1/subdir/myfile myfile
```

- 2) Standard copy; no options. Copy `/pnfs/expt1/subdir/myfile` to a different directory from the cwd, and change the filename:

```
% encp /pnfs/expt1/subdir/myfile \  
/other/dir/newfilename
```

- 3) Request a delayed dismount (**--delayed-dismount**) and get some information printed to screen (**--verbose**). Again, copy `/pnfs/expt1/subdir/myfile` to the cwd, keeping the filename the same.

```
% encp --delayed-dismount=5 --verbose=3 \  
/pnfs/expt1/subdir/myfile myfile
```

- 4) Copy all the files in `/pnfs/expt1/subdir/` starting with the string `trigger1` to the cwd:

```
% encp /pnfs/expt1/subdir/trigger1* .
```

3.2 Encp Command Options

We divide the options into different “risk” categories.

3.2.1 Low risk, anyone can use

The commonly used and “low risk” **encp** user options include:

<code>--data-access-layer</code>	Turns on special status printing; output has standardized format whether error occurred or not.
<code>--delayed-dismout=</code>	Specifies time period in minutes to delay dismount of volume. Use this to tell Enstore: “More work is coming for the volume, don’t dismount the volume too quickly once the current transfer is completed.”
<code>--help</code>	Prints the list of options for encp .
<code>--usage-line</code>	Types out more information about the encp options; not yet implemented.
<code>--verbose=</code>	Changes the amount of information printed about the transfer; provide an integer value. Larger integer numbers provide more “verbosity”. Largest meaningful number may change as development continues.

3.2.2 Know what you’re doing

The second level user options require more knowledge of what you’re doing and more care to be taken. They include:

<code>--age-time=</code>	Specifies the time period, in minutes, after which the priority is eligible to change from the initial job priority. We recommend that you don’t set this, just use the default (60 minutes).
<code>--delpri=</code>	Changes the initial job priority by specified value after a period given by the age-time switch. We recommend that you don’t set this, just use the default (1).
<code>--pnfs-is-automounted</code>	Typically, users should not automount pnfs. If you do, you can specify this option. It alerts the Enstore servers to retry errors due to known OS automounting problems.

Do not use this in non-automounted cases; it can slow the setup of the transfer.

--priority=

Sets the initial job priority to the specified integer value. We recommend that you don't set this, just use the default.



If you feel compelled to set `--priority`, `--delpri` or `--age-time`, please email enstore-admin@fnal.gov first with an explanation, as the defaults should work in almost all cases and changing them may affect other users. Priority goes in strict number sequence, where a higher number means higher priority. Note that Enstore's selection of which file to transfer at a given time uses a much more complicated algorithm than simple priority, however. See Chapter 7: *Job Priority and Queue Management*.

3.3 Copying Files Remotely Using the dCache

If you are off-site and wish to use Enstore, you need to go through a dCache door. You use ftp to copy files back and forth between your machine and your `/pnfs/storage-group` area on the machine running dCache. Enstore and dCache work together (in a manner transparent to the user) to transfer files between storage media and the disk caches on the dCache server. Fermilab's entire `/pnfs/` area is mounted on this server.



The dCache server node and the ports documented in this section are subject to change. You can always find the current ones from the web page http://www-isd.fnal.gov/enstore/dcache_user_guide.html.¹ Currently (February 2002), the dCache server node is `fnlca.fnal.gov`.

3.3.1 Kerberized FTP

The dCache door for Kerberized ftp service enforces Kerberos authentication (see *Strong Authentication at Fermilab Documentation* at <http://www.fnal.gov/docs/strongauth/>). It currently runs on node `fnlca.fnal.gov`, port 24127. Any Kerberized ftp client can be used on the client machine. You must specify the host port in your ftp command, as shown below.

Notes:

1. It is available from the *Fermilab Mass Storage Systems* home page (<http://hppc.fnal.gov/enstore/>); see the list of items under *Documentation* for dCache, and use the *User Access at FNAL* link.

- File reads and writes are supported when the user (a) is authorized by the experiment to access the data stores, and (b) has obtained Kerberos credentials.
- Portal Mode (CRYPTOCARD) access is not supported. It is not compatible with automated transfers or future GRID development.

Sample Kerberized FTP session

User is authenticated to Kerberos and authorized for the Kerberized dCache door (currently at fndca.fnal.gov, port 24127):

```
% ftp fndca.fnal.gov 24127

Connected to stkendca3a.fnal.gov.
220 FTPDoorIM+GSS ready
334 ADAT must follow
GSSAPI accepted as authentication type
GSSAPI authentication succeeded
Name (fndca:aheavey):
200 User aheavey logged in
Remote system type is UNIX.
Using binary mode to transfer files.
ftp> cd aheavey/test3
250 CWD command successful. New CWD is </aheavey/test3>
ftp> ls
200 PORT command successful
150 Opening ASCII data connection for file list
dupl2
duplexps
226 ASCII transfer complete
ftp> get duplexps
local: duplexps remote: duplexps
200 PORT command successful
150 Opening BINARY data connection for /pnfs/fs/usr/test/aheavey/test3/duplexps
226 Closing data connection, transfer successful
42 bytes received in 0.033 seconds (1.2 Kbytes/s)
ftp>
```

3.3.2 Kerberized ftp via the kftpcp Command

A regular ftp client (Kerberized or not) is an interactive program which is hard to use in batch mode. In order to access data from a batch job or a background process, you should either use ftp client libraries (available from many sources), or the **kftp** package, which includes a Kerberized ftp client library for Python.

See section 2.3 *Installing Kftp* for installation information. To use the product in a UPS environment, first run:

```
$ setup kftp
```

Then run the **kftpcp** command to copy one or more files to/from Enstore via the Kerberized dCache door. This command can be used from the shell or in a script.

Syntax and Options

```
% kftpcp [<options>] <source_file> <destination_file>
```

The available options include:

-p <port>	ftp server port number
-m <a p>	ftp server mode; active (default), or passive
-v	verbose mode

Notes:

- If your login id is the same on fndca and your local system, and if they match your Kerberos principal, you can leave off **<your_fndca_login_id>@** in front of **fndca:.**
- Specify the path to the remote file starting from the directory under your **/pnfs/<storage_group>/** area. E.g., to specify the remote file **/pnfs/my_storage_group/path/to/file** on the command line, enter only **/path/to/file**, including the initial slash.
- If you are authorized for multiple storage groups, specify the path to the remote file starting from the storage group you want to use. E.g., to specify the remote file **/pnfs/storage_group_1/path/to/file** on the command line, enter **/storage_group_1/path/to/file**, including the initial slash.

Download a File

To download a stored data file from dCache, run:

```
% kftpcp -p 24127 -m p [-v] \  
[<your_fndca_login_id>@]fndca:</path/to/remote_file> \  
</path/to/local_file>
```

</path/to/remote_file>Upload a File

To upload a new data file, run:

```
% kftpcp -p 24127 -m p [-v] \  
</path/to/local_file> \  
[<your_fndca_login_id>@]fndca:</path/to/remote_file>
```

Examples

To read (download) the stored file **/pnfs/storage_group/mydir/myfile** into a local file of the same name, run:

```
% setup kftp
```

```
% kftpcp -p 24127 -m p -v myloginid@fndca:/mydir/myfile \  
/path/to/myfile
```

Transferred 42 bytes

Or, if your usernames and principal all match, you could run:

```
% kftpcp -p 24127 -m p -v fndca:/mydir/myfile /path/to/myfile
```

Further, if you must specify a particular storage group (**/sg**), run:

```
% kftpcp -p 24127 -m p -v fndca:/sg/mydir/myfile \  
/path/to/myfile
```

3.3.3 NonKerberized FTP (Read-only)

The dCache nonKerberized ftp service currently runs on node fndca.fnal.gov, port 24126. This is read-only, and is not necessarily supported by all experiments. This ftp service can be accessed by ordinary ftp client software. You must specify the host port in your ftp command, as shown below. The Enstore admin will have sent you an email to confirm your registration for this service, and included a password for it.¹ Log in with your username and password.

Sample weakly-authenticated read-only ftp session

Here we explicitly use a nonKerberized ftp client, `/usr/bin/ftp`, and make the connection to fndca port 24126. In the session, we first successfully retrieve a file called `myfile`, and secondly attempt to write a file `trace.txt` and (correctly) fail.

```
% /usr/bin/ftp fndca.fnal.gov 24126  
  
Connected to stkendca3a.fnal.gov.  
220 FTPDoorIM+PWD ready (read-only server)  
Name (fndca:aheavey):  
331 Password required for aheavey.  
Password: (password entered here)  
230 User aheavey logged in  
ftp> cd aheavey/test3  
250 CWD command successful. New CWD is </aheavey/test3>  
ftp> ls  
200 PORT command successful  
150 Opening ASCII data connection for file list  
myfile  
myfile2  
myfile3  
226 ASCII transfer complete  
10 bytes received in 0.018 seconds (0.55 Kbytes/s)  
ftp> get myfile  
200 PORT command successful  
150 Opening BINARY data connection for  
/pnfs/fs/usr/test/aheavey/test3/myfile  
226 Closing data connection, transfer successful
```

1. If you need to change this password, send email to enstore-admin@fnal.gov.


```
local: myfile remote: myfile
42 bytes received in 0.05 seconds (0.82 Kbytes/s)
ftp> put trace.txt
200 PORT command successful
500 Command disabled
ftp> bye
```

